American Economic Journal: Applied Economics 3 (October 2011): 152–185
http://www.aeaweb.org/articles.php?doi=10.1257/app.3.4.152

Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US†

By Olivier Deschênes and Michael Greenstone*

Using random year-to-year variation in temperature, we document the relationship between daily temperatures and annual mortality rates and daily temperatures and annual residential energy consumption. Both relationships exhibit nonlinearities, with significant increases at the extremes of the temperature distribution. The application of these results to "business as usual" climate predictions indicates that by the end of the century climate change will lead to increases of 3 percent in the age-adjusted mortality rate and 11 percent in annual residential energy consumption. These estimates likely overstate the long-run costs, because climate change will unfold gradually allowing individuals to engage in a wider set of adaptations. (*JEL* I12, Q41, Q54)

The climate is a key ingredient in the earth's complex system that sustains human life and well-being. There is a growing consensus that emissions of greenhouse gases due to human activity will alter the earth's climate, most notably by causing temperatures, precipitation levels, and weather variability to increase (Intergovernmental Panel on Climate Change (IPCC) 2007). The development of rational policies requires estimates of the costs associated with these changes in our planet.

Integrated Assessment Models (IAMs) are a popular method to model the costs of climate change. IAMs combine general circulation models of climate and computable general equilibrium economic models to determine the interrelationship between climate and economic activity and policies that affect both of them. An appealing feature of these models is that they allow for a wide range of adaptations

* Deschênes: Department of Economics, University of California, Santa Barbara, 2127 North Hall, Santa Barbara, CA 93106 and Institute for the Study of Labor (IZA) and National Bureau of Economic Research (NBER) (e-mail: olivier@econ.ucsb.edu); Greenstone: Massachusetts Institute of Technology, Department of Economics, E52-359, 50 Memorial Drive, Cambridge, MA 02142-1347 and Brookings Institution and NBER (e-mail: mgreenst@mit.edu). We thank the late David Bradford for initiating a conversation that motivated this project. We thank Maximilian Auffhammer, David Card, Hank Farber, Ted Gayer, Jon Guryan, Solomon Hsiang, Charles Kolstad, David Lee, Erin Mansur, Enrico Moretti, Jesse Rothstein, Bernard Salanie, Matthew White, and Catherine Wolfram for especially valuable criticisms and suggestions. We are also grateful for comments from seminar participants at several universities and conferences. We thank Peter Evangelakis, Elizabeth Greenwood, Ben Hansen, and Henry Swift for exemplary research assistance. We are very grateful to the British Atmospheric Data Centre for providing access to the Hadley Centre data. Greenstone acknowledges the University of California Energy Institute and the Center for Labor Economics at the University of Calfornia, Berkeley for hospitality and support while working on this paper.

† To comment on this article in the online discussion forum, or to view additional materials, visit the article page at http://www.aeaweb.org/articles.php?doi=10.1257/app.3.4.152.

by agents in response to changes in relative prices and climate. Additionally, and perhaps even more importantly, they produce an overall estimate of the impact of climate change or climate change mitigation policies on welfare.

The weakness with IAMs is that inside their black box they can rely on hundreds of parameter estimates. Consequently, the source of the results is unclear, and any notion of uncertainty requires treating many or most of these parameters as known constants. In addition, there is substantial model uncertainty, which is difficult to quantify and rarely reflected in the estimates. Indeed, the Stern Report writes that IAMs are "a computationally demanding exercise, with the result that such models make drastic, often heroic, simplification along all stages of the climate-change chain. What is more, large uncertainties are associated with each element in the cycle" (Stern 2007, 145).[1]

This paper takes a radically different approach than the one used by IAMs to learn about the costs that climate change will impose. Its starting point is that massive historical datasets and transparent statistical models can complement IAMs and even help to improve them. According to the IPCC, human health is a primary channel through which climate change will affect human welfare, yet little is known about the likely impacts (IPCC 2007; World Health Organization (WHO) 2003). Consequently, our focus in this paper is to develop new estimates of the impact of climate change on human health.[2]

The paper's conceptual framework is derived from the canonical Becker-Grossman models of health production (Michael Grossman 2000). Specifically, we estimate the impacts of climate change on mortality *and* expenditures on self-protection or adaptation to learn about the impacts of climate change. Our measure of self-protection is energy consumption, which via air conditioning is perhaps the primary form of protection against high temperatures. The identification strategy relies on the unpredictable and presumably random year-to-year local variation in temperature, so concerns about omitted variables bias are unlikely to be important.

We find a statistically significant relationship between mortality and daily temperatures in the United States, where extremely cold and hot days are associated with elevated mortality rates. However, the magnitude of these effects is smaller than what is suggested by the conventional wisdom based on high profile heat waves (see, e.g., Eric Klinenberg 2002). For example, we find that an additional *day* with a mean temperature exceeding 90° F, relative to a day in the 50°–60° F range, leads to an increase in the *annual* age-adjusted mortality rate of about 0.11 percent. Similarly, a day with a mean temperature below 20° F is associated with an increase in annual mortality of roughly 0.07–0.08 percent. The finding about temperatures and mortality is robust to a wide variety of specification tests and is more pronounced for infants and the elderly. Additionally, there is substantial heterogeneity

[1] In addition, the Stern (2007) Review states "Distributional judgments, a concern with living standards beyond those elements reflected in GDP, and modern approaches to uncertainty all suggest that the appropriate estimate of damages may well lie in the upper part of the range 5–20 percent."

[2] Studies with similar approaches have also appeared in the epidemiology literature, although these studies typically focus on narrower geographical scope (e.g., a subset of US cities) and ignore the role of defensive behavior. See "A Human Health Perspective On Climate Change," Christopher J. Portier et al. (2010) for an overview.

in the responses to extreme temperatures across the country, although it appears unrelated to baseline temperatures or climate.

We also find a nonlinear relationship between annual energy consumption and daily temperature, where energy consumption is elevated in response to temperature-days at the two extremes of the distribution. In particular, there is a proportionally larger increase in energy usage on days where the temperature exceeds 90° F; one extra such *day*, relative to a day in the 50°–60° F range, leads to a 0.4 percent increase in *annual* consumption. The comparable estimate for a day below 20° F is an increase of roughly 0.2 percent–0.3 percent. There is geographic heterogeneity in these responses to extreme temperatures but it appears unrelated to current climate. However, it seems plausible to conclude that the weaker mortality-temperature relationship is at least partially due to the self-protection provided by the cooling from increased energy consumption.

We then combine these estimated relationships with predicted changes in climate from state of the art climate models and "business as usual" scenarios to develop estimates of the health related welfare costs of climate change in the United States. The preferred estimates suggest that climate change will lead to an increase in the age adjusted US mortality rate of 3 percent by the end of the twenty-first century. To put this in context, the age-adjusted mortality rate has declined by roughly 1 percent per year over the last 35 years, so this appears to be a small effect. In contrast, the energy results suggest that climate change will cause annual US residential energy consumption to increase by up to 11 percent at the end of the century. Taken together, the present discounted values of the losses due to increased mortality and energy consumption over the remainder of the century range from 10 percent to 58 percent of 2010 GDP, but the null hypothesis of a zero loss cannot be rejected at conventional significance levels.

The analysis is conducted with detailed and comprehensive data available on mortality, energy consumption, weather, and climate change predictions for fine US geographic units covering the whole continental United States. The mortality data come from more than 73.2 million death certificates in the 1968–2002 Compressed Mortality Files, the energy data are from the Energy Information Administration and the weather data are from the thousands of weather stations located throughout the United States. We utilize two sets of state of the art end-of-century (i.e., 2070–2099) daily climate change predictions that represent "business-as-usual" scenarios.

Finally, the paper's approach mitigates or solves the conceptual and statistical problems that have plagued previous research (e.g., Francesco Bosello, Roberto Roson, and Richard S. J. Tol 2006, Katharine Hayhoe et al. 2004). First, the availability of data on energy consumption means that we can measure the impact on mortality *and* self-protection expenditures. Second, the estimation of annual mortality equations, rather than daily ones, mitigates concerns about failing to capture the full mortality impacts of temperature shocks. Third, the use of comprehensive datasets from the entire United States over a 35 year period means that the estimates are externally valid to a wide range of places. Specifically, they are more general than estimates derived from studies of single events in small geographic areas over a small fraction of the temperature distribution (e.g., the impact of the French heat wave) that have been emphasized by the IPCC and Stern Report (IPCC 2007; Stern

2007). Fourth, the statistical models include county and state-by-year fixed effects so they adjust for any permanent differences in unobserved health across counties due to sorting while also controlling for differential trends across states. Fifth, we model daily temperature semi-parametrically, so we do not rely on functional form assumptions to infer the impacts of the hottest and coldest days. Finally, we estimate separate mortality models for four age groups, which allows for substantial heterogeneity in the impacts of temperature.

Before proceeding, we emphasize that there are a few important caveats to our effort to estimate the health related welfare impacts of climate change-induced increases in temperature. On the one hand, the estimated impacts likely overstate the mortality and adaptation costs, because the analysis relies on inter-annual variation in weather, and less expensive adaptations will be available in response to permanent climate change. On the other hand, the estimated welfare losses fail to include the impacts on other health-related determinants of welfare (e.g., morbidities and diseases vectors) that may be affected by climate change, so in this sense they are an underestimate. Additionally, the effort to project outcomes at the end of the century requires a number of strong assumptions, including that the climate change predictions are correct, relative prices (e.g., for energy and medical services) will remain constant, the same energy and medical technologies will prevail, and the demographics of the US population (e.g., age structure) and their geographic distribution will remain unchanged. These are strong assumptions, but in contrast to IAMs they allow for a transparent analysis based on data rather than on unverifiable assumptions.

The paper proceeds as follows. Section I briefly reviews the patho-physiological and statistical evidence on the relationship between weather and mortality. Section II provides the conceptual framework for our approach. Section III describes the data sources and reports summary statistics. Section IV presents the econometric approach, and Section V describes the results. Section VI assesses the magnitude of our estimates of the effect of climate change and discusses a number of important caveats to the analysis. Section VII concludes the paper.

I. Background on the Relationship between Weather and Mortality

Individuals' heat regulation systems enable them to cope with high and low temperatures. Specifically, high and low temperatures generally trigger an increase in the heart rate in order to increase blood flow from the body to the skin, leading to the common thermoregulatory responses of sweating in hot temperatures and shivering in cold temperatures. These responses allow individuals to pursue physical and mental activities without endangering their health within certain ranges of temperature. Temperatures outside of these ranges pose dangers to human health and have been empirically shown to result in premature mortality.[3]

The challenge for this study, and for any study focused on substantive changes in life expectancy, is to develop estimates of the impact of temperature on mortality that are based on the full long-run impact on life expectancy. In the case of hot days,

[3] See Rupa Basu and Jonathan M. Samet (2002) and Deschênes and Greenstone (2008) for reviews of this literature.

the previous literature suggests that this task requires purging the temperature effects of the influence of harvesting or forward displacement. In the case of cold days, the mortality impact may accumulate over time. In both cases, the key point is that the full impact of a given day's temperature may take numerous days to manifest fully.

Our review of the literature suggests that the full mortality impacts of cold and hot days are likely to be concentrated within 30 days of the exposure (Maud M. T. E. Huynen et al. 2001; Deschênes and Enrico Moretti 2009). The econometrics section below outlines a method that allows the mortality impacts of temperature to manifest over long periods of time. Further, the immediate and longer run effects of hot and cold days are likely to vary across the populations, with larger impacts among relatively unhealthy subpopulations. One important determinant of health is age, with the old and young being particularly sensitive to environmental stresses. Consequently, we conduct separate analyses for four separate age categories.

II. Conceptual Framework

The paper's goal is to develop a partial estimate of the health-related welfare impact of temperature changes in the United States. This section begins by reviewing a Becker-Grossman style 1-period model of health production (Grossman 2000). It then uses the results to derive a practical expression for the health-related welfare impacts of temperature changes (Winston Harrington and Paul R. Portney 1987). This expression guides the subsequent empirical analysis. The section then discusses the implications of our estimation strategy, which relies on inter-annual fluctuations in weather to develop an estimate of the welfare impacts of the health-related welfare impacts of temperature increases due to climate change.

A Practical Expression for Willingness to Pay/Accept (WTP/WTA) for an Increase in Temperature.—We assume a representative individual consumes a jointly aggregated consumption good, x_C. Their other consumption good is the probability of survival, which leads to a utility function of

$$U = U[x_C, s], \tag{1}$$

where s is the survival rate. The production function for survival is expressed as

$$s = s(x_H, T), \tag{2}$$

so that survival is a function of x_H, which is a private good that increases the probability of survival, and ambient temperature, T. Energy consumption is an example of x_H, since energy is used to run air conditioners, which affect survival on hot days. We define x_H such that $\partial s/\partial x_H > 0$. For this section's expositional purposes, we consider temperature increases during the summer only when higher temperatures are harmful for health so $\partial s/\partial T < 0$.

The individual faces a budget constraint of the form:

$$I - x_C - px_H = 0, \tag{3}$$

where I is exogenous earnings or income, and prices of x_C and x_H are 1 and p, respectively.

The individual's problem is to maximize (1) through her choices of x_C and x_H, subject to (2) and (3). In equilibrium, the ratio of the marginal utilities of consumption of the two must be equal to the ratio of the prices: $[(\partial U/\partial s) \cdot (\partial s/\partial x_H)]/[\partial U/\partial x_C] = p$. Solution of the maximization problem reveals that the input demand equations for x_C and x_H are functions of prices, income, and temperature. Further, it reveals the indirect utility function, V, which is the maximum utility obtainable given p, I, and T.

We utilize $V(p, I, T)$ to derive an expression for the welfare impact of temperature increases, holding constant utility (and prices). In this case, it is evident that the consumer must be compensated for changes in T with changes in I when utility is held constant. The point is that in this setting income is a function of T, which we denote as $I^*(T)$. Consequently, for a given level of utility and fixed p, there is an associated $V(I^*(T), T)$.

Now, consider the total derivative of V with respect to T along an indifference curve:

$$dV/dT = V_I \cdot \left(dI^*(T)/dT\right) + \partial V/\partial T = 0 \qquad \text{or}$$

$$dI^*(T)/dT = -(\partial V/\partial T)/(\partial V/\partial I).$$

The term $dI^*(T)/dT$ is the change in income necessary to hold utility constant for a change in T. In other words, it measures willingness to pay (accept) for a decrease (increase) in summer temperatures.

Since the indirect utility function is not observable, it is useful to express $dI^*(T)/dT$ in terms that can be measured with available datasets. By using the derivatives of V and the first order conditions from the above maximization problem, it can be rewritten as $dI^*(T)/dT = -p\ [(\partial s/\partial T)/(\partial s/\partial x_H)]$. In principle, it is possible to measure these partial derivatives, but it is likely infeasible since data files containing measures of the complete set of x_H are unavailable generally. Put another way, data limitations prevent the estimation of the production function specified in equation (2). However, a few algebraic manipulations, based on the first order conditions and that $\partial s/\partial T = ds/dT - (\partial s/\partial x_H)(\partial x_H/\partial T)$ (because $ds/dT = (\partial s/\partial x_H)(\partial x_H/\partial T) + \partial s/\partial T$), yield:

$$dI^*(T)/dT = -ds/dT\,(\partial U/\partial s)/\lambda + p\,\partial x_H/\partial T, \tag{4}$$

where λ is the Lagrangian multiplier from the maximization problem or the marginal utility of income.

As equation (4) makes apparent, willingness to pay/accept for a change in temperature can be inferred from changes in s and x_H. Since temperature increases raise the effective price of survival, theory would predict that $ds/dT \leq 0$ and $\partial x_H/\partial T \geq 0$. Depending on the exogenous factors, it is possible that there will be a large change in the consumption of x_H (at the expense of consumption of x_C) and little change in s. The key point for this paper's purpose is that the full welfare effect of the exogenous change in temperature is reflected in changes in the survival rate *and* the consumption of x_H.

It is of tremendous practical value that all of the components of equation (4) can be measured. The total derivative of the survival function with respect to temperature (ds/dT), or the dose-response function, is obtained through the estimation of epidemiological-style equations that do not control for x_H. We estimate such an equation below.[4] The term $(\partial U/\partial s)/\lambda$ is the dollar value of the disutility of a change in the survival rate. This is known as the value of a statistical life (Richard Thaler and Sherwin Rosen 1976) and empirical estimates are available (e.g., Orley Ashenfelter and Greenstone 2004). The last term is the partial derivative of x_H with respect to temperature multiplied by the price of x_H. We estimate how energy consumption changes with temperature (i.e., $\partial x_H/\partial T$) below, and information on energy prices is readily available.

It is appealing that the paper's empirical strategy can be directly connected to an expression for WTP/WTA, but this connection has some limitations worth highlighting from the outset. The empirical estimates will only be a partial measure of the health-related welfare loss, because temperature changes may affect other health outcomes (e.g., morbidity rates). Further, although energy consumption likely captures a substantial component of health preserving or defensive expenditures, temperature changes may induce other forms of adaptation (e.g., substituting indoor exercise for outdoor exercise or changing the time of day when one is outside).[5] These other outcomes are unobservable in our data files, so the resulting welfare estimates will be incomplete and understate the costs of the temperature shocks observed in our data.

The Welfare Impacts of Temperature Shocks and Climate Change.—It is intuitive to apply the one-period model sketched in the previous subsection to an examination of the welfare consequences of inter-annual fluctuations in weather. For example, it is easy to turn the thermostat down and use more air conditioning on hot days, and it is even possible to purchase an air conditioner in response to a single year's heat wave. However, this simple model fails to capture the full set of adaptations that households would undertake in response to permanent climate change. For instance, permanent climate change is likely to lead individuals to make their homes more energy efficient or perhaps even to migrate (presumably to the North). These longer run investments or adaptations are not captured in this model and are also missing from the empirical results due to the empirical strategy that relies on inter-annual variation in weather. The point is that the paper's estimates can be used to develop credible estimates of the welfare costs of inter-annual variation in weather, but these estimates will overstate the measurable health-related welfare costs.

[4] Previous research on the health impacts of air pollution almost exclusively estimate these dose-response functions, rather than the production functions specified in equation (2) (e.g., Kenneth Y. Chay and Greenstone 2003).

[5] Energy consumption may affect utility through other channels in addition to its role in self-protection. For example, high temperatures are uncomfortable. It would be straightforward to add comfort to the utility function and make comfort a function of temperature and energy consumption. In this case, this paper's empirical exercise would fail to capture the impact of temperature on heat, but the observed change in energy consumption would reflect its role in self-protection *and* comfort.

III. Data Sources and Summary Statistics

To implement the analysis, we collected the most detailed and comprehensive data available on mortality, energy consumption, weather, and predicted climate change. This section briefly describes these data and reports summary statistics. More details on the data sources are provided in the online Appendix.

A. *Data Sources*

Mortality and Population Data.—The mortality and population data are taken from the Compressed Mortality Files (CMF) compiled by the National Center for Health Statistics.[6] The CMF contains the universe of the 72.3 million deaths in the United States from 1968 to 2002. Importantly, the CMF reports death counts by race, sex, age group, county of residence, cause of death, and year of death. In addition, the CMF files also contain population totals for four age groups, which we use to calculate all-cause and cause-specific mortality rates. We complemented these data with the 1972 to 1988 Multiple Cause of Death files. Our initial sample consists of all deaths occurring in the continental 48 states plus the District of Columbia.

Energy Data.—The energy consumption data comes directly from the U.S. Energy Information Administration (EIA) State Energy Data System (EIA 2006). These data provide state-level information about energy prices, expenditures, and consumption from 1968 to 2002. The data is disaggregated by energy source and end use sector. All energy data is given in British Thermal Units (BTU).

We used the database to create an annual state-level panel data file for total energy consumption by the residential sector, which is defined as "living quarters for private households." The database also reports on energy consumption by the commercial, industrial, and transportation sectors. These sectors are not a focus of the analysis, because they do not map well into the health production function model outlined in Section II. Further, factors besides temperature are likely to be the primary determinant of consumption in these sectors.

The measure of total residential energy consumption is comprised of two pieces: "primary" consumption, which is the actual energy consumed by households, and "electrical system energy losses." The latter accounts for about two-thirds of total residential energy consumption. It is largely due to losses in the conversion of heat energy into mechanical energy to turn electric generators, but transmission and distribution and the operation of plants also account for part of the loss. In the 1968–2002 period, total residential energy consumption increased from 7.3 quadrillion (quads) BTU to 21.2 quads, and the mean over the entire period was 16.6 quads.

Weather Data.—The weather data are drawn from the National Climatic Data Center (NCDC) Summary of the Day Data (File TD-3200). The key variables for our analysis are the daily maximum and minimum temperature as well as the total

[6] From 1989 onward, the CMF are not in the public domain. We gained access to these data through an agreement with the Centers for Disease Control.

daily precipitation.[7] For most parts of the paper, we follow the convention in the literature and focus on the daily mean temperature, which is the simple average of the maximum and minimum.

To ensure the accuracy of the weather readings, we developed a rule to select the weather stations that requires monitors to operate for a minimum number of days (see the online Data Appendix for details). The acceptable station-level data is then aggregated at the county level by taking an inverse-distance weighted average of all the valid measurements from stations that are located within a 200 km radius of each county's centroid, where the weights are the inverse of their squared distance to the centroid so that more distant stations are given less weight. This procedure yields a balanced panel of 3,066 counties with acceptable weather data that accounts for 99 percent of all deaths in the United States from 1968 to 2002.

Climate Change Prediction Data.—Climate predictions are based on two state of the art global climate models. The first is the Hadley Centre's third Coupled Ocean-Atmosphere General Circulation Model, which we refer to as Hadley 3. This is the most complex and recent model in use by the Hadley Centre. We also use predictions from the National Center for Atmospheric Research's (NCAR) Community Climate System Model (CCSM) 3, which is another coupled atmospheric-ocean general circulation model (NCAR 2007). Both models are considered state of the art and were used in the fourth Assessment Report by the IPCC (IPCC 2007).

Predictions of climate change from both of these models are available for several emission scenarios, corresponding to "storylines" describing the way the world (population, economies, etc.) may develop over the next 100 years. We focus on the A1FI and A2 scenarios. These are "business-as-usual" scenarios, which are the proper scenarios to consider when judging policies to restrict greenhouse gas emissions.

We obtained *daily* temperature and precipitation predictions for grid points throughout the continental United States from the application of A1FI scenario to the Hadley 3 model for the years 1990–2099 and the A2 scenario to the CCSM 3 for the years 2000–2099. The Hadley model gives daily minimum and maximum temperatures, while the CCSM model only reports the average of the minimum and maximum. Each set of predictions is based on a single run of the relevant model and available for an equidistant set of grid points over land in the United States.

We calculate future temperature and precipitation realization in two ways. The first assigns each county a daily weather realization directly from the Hadley and CCSM predictions. Specifically, this is calculated as the inverse-distance weighted average among all grid points within a given distance from the county's centroid.

[7] Other aspects of daily weather, such as humidity and wind speed, could influence mortality, both individually and in conjunction with temperature. Importantly for our purposes, there is little evidence that wind chill factors (a nonlinear combination of temperature and wind speed) perform better than simple temperature levels in explaining daily mortality rates (Anton E. Kunst, Feikje Groenhof, and Johan P. Mackenbach 1994). In addition, there are relatively few weather stations that measure relative humidity. For example, only 308 stations measured relative humidity in July 2000. The corresponding figures for July 1990 and 1980 are 259 and 0, respectively. As such, the coverage of stations that record humidity is too sparse for our analysis. Alan Barreca (2009) examines the relationship between humidity and mortality in the United States.

The limitation of this approach is that there may be systematic model error that causes the predictions to be too high or too low.

The second method adjusts the Hadley 3 A1FI weather predictions for model error by comparing the model's predictions for the 1990–2002 period with the actual realizations from the weather station data. For example, in the case of temperature, we calculate Hadley 3 model errors for each of the 365 days in a year separately for each county as the average difference between county by day of year specific average temperature from the weather station data and the Hadley 3 A1FI predictions during the 1990–2002 period. This county by day of year-specific error is then added to the Hadley 3 A1FI predictions to obtain an error-corrected climate change prediction. The main limitations of this approach are that the historical Hadley predictions are only available for some years (i.e., 1990–2002) that were hot by historical standards and thirteen years is a relatively short period to validate the model.

In the subsequent analysis, we focus on the Hadley 3 A1FI error-corrected predictions. However, we will also report estimates based on the Hadley 3 A1FI and CCSM 3 A2 predictions that are not corrected for model error in some of the tables. Again, it is impossible to derive error-corrected CCSM 3 A2, because historical predictions are unavailable.

B. *Summary Statistics*

Weather and Climate Change Statistics.—The analysis uses the rich daily historical weather and daily climate change predictions data to develop county-level measures of past and future weather. Table 1 reports on national and regional measures of observed temperature from 1968 to 2002 (reported in panel A) and predicted changes, based on the difference between the 2070–2099 prediction and the 1968–2002 baseline (reported in panel B). For the historical temperature, this is calculated across all county-by-year observations and weighted by the total county population in the year. The predicted temperature under climate change is based on the average of 2070–2099 predictions, where the weight is the average total county population over the years 1968–2002. Since these calculations of actual and predicted temperatures depend on the geographic distribution of the population in the United States, systematic migration (e.g., from South to North) would change them without any change in the underlying climate.

In each panel, the first row reports averages for all counties, while the following nine columns are for the nine US census divisions. The average daily mean temperature in column 1 is 56.4° F. This reflects the variation across all years and counties, as well as the within-year variation. The entries for the nine census divisions reveal the geographical variation in this average ranges from 48.5° F (New England) to 65.7° F (West South Central). Nationally, the average difference between the minimum and maximum is 21° F.

The pale bars in Figure 1 depict the average annual distribution of daily mean temperatures across ten temperature categories or bins, again, during the 1968–2002 period. These categories represent daily mean temperature less than 10° F, greater than 90° F, and the eight 10° F wide bins in between. The height of each light-colored bar corresponds to the mean number of days that the average person experiences in

Table 1—Historical Averages of Daily Weather Across Counties, and Predicted Changes in Error-Corrected Hadley 3 A1FI

	Average daily temperature			Days with mean > 90° F		Days with mean < 10° F	
	Mean (1)	Minimum (2)	Maximum (3)	Number of days per year (4)	Average daily temperature (5)	Number of days per year (6)	Average daily temperature (7)
Panel A. Historical temperature data (1968–2002)							
All counties	56.4	45.9	66.9	1.39	2.7	3.9	3.2
By US census division							
1. New England	48.5	38.5	58.6	0.0	90.8	6.1	5.0
2. Middle Atlantic	51.8	42.4	61.3	0.1	91.0	3.0	5.3
3. East North Central	49.8	40.1	59.6	0.09	0.9	8.6	3.9
4. West North Central	50.2	39.4	61.0	0.59	1.5	16.1	1.1
5. South Atlantic	62.8	52.3	73.3	0.1	91.1	0.3	6.0
6. East South Central	59.9	48.8	71.0	0.19	1.0	0.7	5.1
7. West South Central	65.7	54.8	76.7	2.69	1.4	0.2	6.6
8. Mountain	56.1	42.2	70.0	13.7	93.1	3.6	2.3
9. Pacific	59.9	48.9	70.8	1.19	2.3	0.2	4.1
Panel B. Predicted change, error-corrected Hadley 3 A1FI, (2070–2099)							
All counties	11.2	10.7	11.9	42.3	3.1	−3.4	1.4
By US census division							
1. New England	11.6	11.4	11.8	21.4	4.4	−5.8	1.6
2. Middle Atlantic	12.1	11.9	12.3	34.9	5.3	−3.0	2.0
3. East North Central	11.7	11.2	12.1	33.8	5.3	−7.5	0.8
4. West North Central	12.3	11.8	12.8	47.0	4.9	−13.2	3.0
5. South Atlantic	10.1	9.9	10.4	43.8	4.1	−0.3	2.5
6. East South Central	10.1	9.8	10.5	55.5	4.4	−0.7	1.7
7. West South Central	10.3	10.2	10.4	95.9	4.4	−0.2	0.7
8. Mountain	10.5	10.1	11.0	28.6	4.0	−3.2	2.1
9. Pacific	9.0	7.9	10.1	23.9	2.1	−0.1	0.4

Notes: Averages are calculated for samples size 39,168,150 (panel A) and 33,572,700 (panel B). These represent the 365 annual observations for the 3,066 counties in the sample, over a 35 year period (panel A) or a 30 year period (panel B). Averages are weighted by total county population in a given year (panel A) or the average total county population during 1968–2002 (panel B). Predicted change defined as difference between 2070–2099 average and 1968–2002 average. Census division defined as follows: New England (CT, ME, MA, NH, RI, VT), Middle Atlantic (NJ, NY, PA), East North Central (IN, IL, MI, OH, WI), West North Central (IA, KS, MN, MO, NE, ND, SD), South Atlantic (DE, DC, FL, GA, MD, NC, SC, VA, WV), East South Central (AL, KY, MS, TN), West South Central (AR, LA, OK, TX), Mountain (AZ, CO, ID, NM, MT, UT, NV, WY), and Pacific (CA, OR, WA). See the text for more details.

each bin. This is calculated as the weighted average across county-by-year realizations, where the county-by-year's total population is the weight.[8] The average number of days in the modal bin of 60°–70° F is 74.6. The mean number of days at the endpoints is 3.9 for the less than 10° F bin and 1.3 for the greater than 90° F bin. These ten bins form the basis for our semi-parametric modeling of temperature in equations for mortality rates and energy consumption throughout the remainder of the paper. This binning of the data preserves the daily variation in temperatures,

[8] For a given county-by-year, the number of days in each bin is calculated as the inverse-distance weighted average of the number of days in each bin at all weather stations within a 200 kilometer radius of the county centroid. This preserves the variation in temperatures, relative to assigning days to bins after averaging across mean temperatures from weather stations.

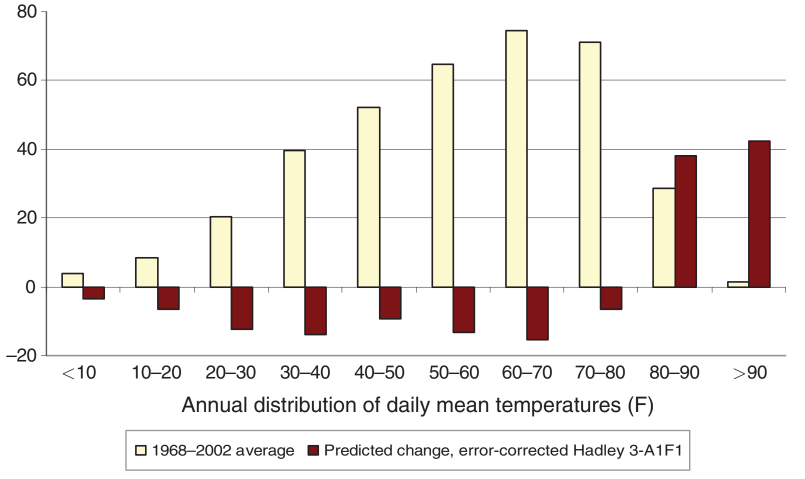

FIGURE 1. DISTRIBUTION OF ANNUAL DAILY MEAN TEMPERATURES (F), 1968–2002, AND PREDICTED CHANGES IN DISTRIBUTION ACCORDING TO ERROR-CORRECTED HADLEY 3 A1FI

Notes: Figure 1 shows the historical average and the predicted change in the distribution of daily mean temperatures across ten temperature-day bins. The "1968–2002 average" bars represent the average number of days per year in each temperature category for the 107,590 county-year observations in the sample, weighted by the total population in a county-year. The "predicted change, error-corrected Hadley 3A1FI" bars represent the change in the average number of days per year in each temperature category. Changes are defined as the difference between the 1968–2002 average in each category and the 2070–2099 predicted average number of days in each category. Predicted changes are weighted by the average total population over 1968–2002 in a county. See the text for more details.

which is an improvement over much of the previous research that obscures much of the variation in temperature.[9] This is important because of the potential for substantial nonlinearities in the daily temperature-mortality and daily temperature-energy consumption relationships.

Returning to Table 1, columns 4–7 in panel A report average temperatures for days when the mean exceeds 90°F or is less than 10°F. These correspond to the two extremes of the temperature distribution, where the mortality impacts are likely to be more pronounced. Exposure to extreme hot temperatures is infrequent as the national average is 1.3 days per year, and this highlights that identifying the effect of exposure to extreme heat on mortality is empirically challenging. Moreover there are significant geographical differences in exposure, which ranges from 0.01 days per year to 13.7 days per year in the New England and mountain divisions, respectively. As such, most of the identification of this effect will come from southern and western states. Columns 6 and 7 report statistics on the frequency of extreme cold days, where the mean temperature is 10°F or less. Across all counties the average is 3.9 days per year, and there is a great degree of variation across census divisions. Again, the regional variation is notable and most of the identification of the effect

[9] For example, W. J. M. Martens (1998) and Tol (2002) use the maximum and the minimum of monthly mean temperatures over the course of the year.

of cold exposure on mortality will come from northern states (New England, East North Central, and West North Central).

Figure 1 also provides an opportunity to understand how the full distributions of daily mean temperatures are expected to change. The dark bars report the change in the number of days in each temperature category that the average person is expected to experience. The most important changes in the distribution are in the last two bins. The error-corrected Hadley 3 A1FI predictions indicate that the typical person will experience 38.0 *additional* days per year where the mean daily temperature is between 80° F and 90° F. Equally important is that the mean daily temperature is predicted to exceed 90° F for 42.3 *additional* days per year.[10] To put this in perspective, the average person currently experiences 28.7 days in the 80°–90° F range and just 1.3 days per year where the mean exceeds 90° F.[11] An examination of the remainder of the figure highlights that the reduction in extreme cold days is much smaller than the increase in extreme hot days. The subsequent analysis demonstrates that this has a profound effect on the estimated impacts of climate change on mortality and energy consumption.

Panel B in Table 1 reports the predicted change in each of the temperatures. The predicted change is based on the departure of the 2070–2099 average from the 1968–2002 baseline.[12] The error-corrected Hadley 3 A1FI predictions indicate a change in mean temperature of 11.2° F. The remaining rows of panel B document the extent of geographical heterogeneity in the most extreme daily temperature days (less than 10° F and more than 90° F).

IV. Econometric Strategy

This section describes the econometric models for annual mortality rates and residential energy consumption.

Mortality Rates.—We fit the following equation:

$$(5)\quad Y_{cta} = \sum_{j} \theta_{aj}^{TMEAN}\, TMEAN_{ctj} + \sum_{l} \delta_{al}^{PREC}\, PREC_{ctl} + \alpha_{ca} + \gamma_{sta} + \varepsilon_{cta},$$

where Y_{cta} is the mortality rate for age group a in county c in year t. In the subsequent analysis, this equation is estimated separately for four separate age groups (ages 0–1, 1–44, 45–64, and 65+), so that all parameters are allowed to vary across these age groups. The last term in equation (5) is the stochastic error term, ε_{cta}.

The variables of interest are the measures of temperature and precipitation. They are constructed to capture the full distribution of annual fluctuations in weather.

[10] We emphasize that a mean daily temperature of 90° F is very hot. For example, a day with a high of 100° F would need a minimum temperature greater than 80° F to qualify.

[11] The paper's primary econometric model for mortality assumes that the impact of all days in the greater than 90° F bin have an equal impact. As Table 1 shows, the predicted increase in mean temperatures among days in this bin is relatively modest, with predicted increases of 3.1° F from the error-corrected Hadley 3 A1FI. Consequently, historical data are likely to be informative about the impacts of the additional days with temperatures exceeding 90° F.

[12] For comparability, we follow much of the previous literature on climate change and focus on the temperatures predicted to prevail at the end of the century.

The variables $TMEAN_{ctj}$ denote the number of days in county c and year t where the daily mean temperature is in the jth of the ten bins used in Figure 1. Thus, the only functional form restriction is that the impact of the daily mean temperature on the annual mortality rate is constant within 10° F intervals. The choice of ten temperature bins represents an effort to allow the data, rather than parametric assumptions, to determine the mortality-temperature relationship, while also obtaining estimates that are precise enough that they have empirical content. This degree of flexibility and freedom from parametric assumptions is only feasible because we are using 35 years of data from the entire United States. The variables $PREC_{ctl}$ are simple indicator variables based on annual rainfall in county c in year t. Each indicator corresponds to a 5-inch bin and there are 11 of them, ranging from less than 10 inches to more than 60 inches.

The equation includes a full set of county-by-age group fixed effects, α_{ca}, which absorb all unobserved county-specific time invariant determinants of the mortality rate for each age group. So, for example, differences in permanent hospital quality or the overall healthiness of the local age-specific population will not confound the weather variables. The equation also includes state-by-year effects, γ_{sta}, that are allowed to vary across the age groups. These fixed effects control for time-varying differences in the dependent variable that are common across counties within age groups in a state (e.g., changes in state Medicare policies).

The validity of this paper's empirical exercise rests crucially on the assumption that the estimation of equation (5) will produce unbiased estimates of the θ_{aj}^{TMEAN} and δ_{al}^{PREC} vectors. By conditioning on the county and state-by-year fixed effects, these vectors are identified from county-specific deviations in weather about the county averages after controlling for shocks common to all counties in a state. Due to the unpredictability of weather fluctuations, it seems reasonable to presume that this variation is orthogonal to unobserved determinants of mortality rates. For example, in the case of the jth temperature variable, we believe the identifying assumption $E[TMEAN_{ctj}\, \varepsilon_{cta} \,|\, TMEAN_{ct-j}, PREC_{ctl}, \alpha_{ca}, \gamma_{sta}] = 0$ is valid.

There are three further issues about the econometric approach that bear noting. First, it is likely that the error terms are correlated within county-by-age groups over time. Consequently, the paper reports standard errors that allow for heteroskedasticity of an unspecified form and that are clustered at the county-by-age group level. Second, we fit weighted versions of equation (5), where the weight is the square root of the age group's population in the county (i.e., the denominator) for two complementary reasons. The estimates of mortality rates from large population counties are more precise, so it corrects for heteroskedasticity associated with these differences in precision. Further, the results reveal the impact on the average person, rather than on the average county, which we believe is more meaningful.

Third, in many places, the paper reports age-adjusted estimates of the impact of temperature on mortality. These are calculated as the weighted averages of the regression coefficients $\left(\text{i.e., } \hat{\theta}_{aj}^{TMEAN}\right)$ from the estimation of the four age group-specific regressions. The national weights are 0.015 for ages 0–1, 0.667 for ages 1–44, 0.200 for 45–64, and 0.117 for ages 65 and older, which represent these age categories' population shares during the sample period, 1968–2002.

Residential Energy Consumption.—We fit the following equation for state-level residential energy consumption:

$$(6)\quad \ln(C_{st}) = \sum_j \theta_j^{TMEAN}\, TMEAN_{stj} + \sum_l \delta_l^{PREC}\, PREC_{stl} + \mathbf{X}_{st}\beta + \alpha_s + \gamma_{dt} + \varepsilon_{st}.$$

C_{st} is residential energy consumption in state s in year t, and d indexes US census divisions. The modeling of temperature and precipitation is identical to the approach in equation (5). The only difference is that these variables are measured at the state-by-year level. They are calculated as the weighted average of the county-level versions of the variables, where the weight is the county's population in the relevant year. The equation also includes state fixed effects (α_s), census division-by-year fixed effects (γ_{dt}), and a stochastic error term, ε_{st}.

The vector $\mathbf{X}_{st}$ includes the state-level ln of population and gross domestic product (GDP) and their squares. The former variable accounts for differences in population growth within a census division that might confound the consistent estimation of the θ_js.[13] The latter is included because energy consumption is a function of income.

Finally, we will also report the results from versions of equation (6) that model temperature with heating and cooling degree days. We follow the conventional approach and use a base of 65°F to calculate both variables. To implement this alternative method for modeling a year's temperature, we sum the number of heating and cooling degree days separately over the year and control for these variables in equation (6) instead of the $TMEAN_{stj}$ variables.

V. Results

This section is divided into four subsections. The first provides estimates of the relationship between daily temperatures, and the second explores the source of the relationship between extreme temperatures and annual mortality rates. The third section reports estimates of the effect of temperature on residential energy consumption. Finally, the last subsection uses these estimated relationships to predict the impacts of climate change on annual mortality and residential energy consumption.

A. *Estimates of the Impact of Temperature on Mortality Based on Annual Data*

Figure 2 presents the preferred estimates of the impact of exposure to temperature on annual mortality from the estimation of equation (5). Specifically, it plots the age-adjusted regression coefficients $\left(\text{i.e., } \hat{\theta}_j^{TMEAN}\right)$. In each of the underlying age-specific regressions, the $TMEAN_j$ variable associated with the 50°–60°F bin was dropped, so each $\hat{\theta}_j^{TMEAN}$ measures the estimated impact of an additional day in bin j on the mortality rate (i.e., deaths per 100,000), relative to the impact of a day in the 50°–60°F range. The figure also plots the estimated $\hat{\theta}_j^{TMEAN}$s plus and minus two standard errors, so their precision is apparent.

[13] For example, Arizona's population has increased by 223 percent between 1968 and 2002 compared to just 124 percent for the other states in its census division.

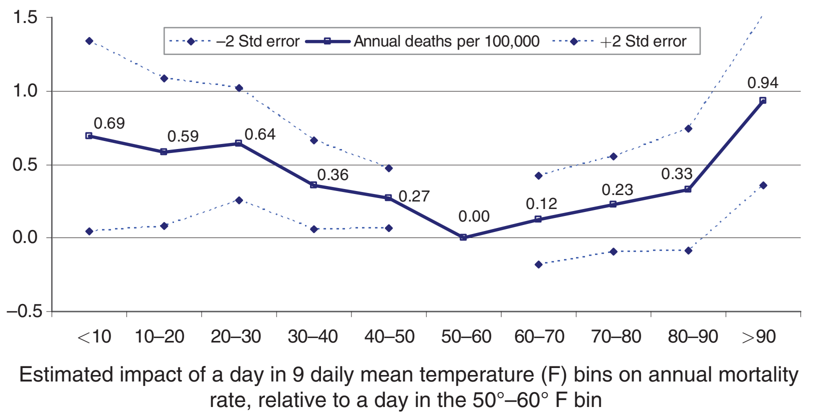

Figure 2. Estimated Relationship Between Annual Age-Adjusted Mortality Rate Per 100,000 and Average Daily Temperature

Notes: Figure 2 plots the aggregate response function between annual mortality rate (per 100,000) and average daily temperatures. This is obtained by fitting equation (5) for the annual mortality rate in each age group. The age group-specific estimates are then combined into a single age-adjusted estimate by taking a weighted average of the age-specific estimates, where the weight is the age group's share of the total population. The response function is normalized with the 50º–60ºF category set equal to 0 so each θ_j corresponds to the estimated impact of an additional day in bin j on the annual age-adjusted mortality rate (i.e., deaths per 100,000) relative to the mortality rate associated with a day where the temperature is between 50º–60ºF. The figure also plots the $\hat{\theta}_j^{TMEAN}$s plus and minus two standard error of the estimates. The numbers above the response function correspond to the point estimates associated with each temperature bin. See the text for more details.

It is evident that mortality risk is highest at the coldest and hottest temperatures. Indeed, the response function between mortality and temperature has a flattened U-shape. For example, the mean of the coefficients associated with the three lowest temperature categories is 0.62, so exchanging a single day in this range for one in the 50°–60° F range would lead to 0.62 fewer annual deaths per 100,000. The largest coefficient is the one associated with the >90° F temperature bin and it is 0.94. It is noteworthy that the null of equality with the base category can be rejected at conventional significance levels for these hot and cold extreme temperatures, even though there are relatively few days to identify the coefficients (recall Table 1). Finally, the coefficients associated with the temperatures in the middle range are all smaller in magnitude, and it is frequently impossible to reject the null of equality with the base category.

Table 2 reports the point estimates associated with the temperature variables from the fitting of several versions of equation (5) to better characterize the results and evaluate their robustness across alternative samples and subpopulations. For compactness, only the point estimates corresponding to the lowest two (coldest) and highest two (hottest) exposure bins are reported, although all models include the full set of nine temperature bin variables. Panel A reports the baseline estimates separately by age as well as the age-adjusted estimates that were used to create

Table 2—Estimates of the Impact of Extreme Temperatures on Annual Mortality Rate, By Age Group and Across Alternative Specifications

	Impact on annual mortality rate per 100,000			
	Days < 10° F (1a)	Days 10°–20° F (1b)	Days 80°–90° F (1c)	Days > 90° F (1d)
Panel A. Baseline estimates, all cause-mortality				
1. Infant (0–1)	−0.637	−0.348	0.736	2.056
	(1.036)	(0.798)	(0.414)	(1.091)
2. Age 1–44	0.178	0.278*	0.193*	0.224*
	(0.111)	(0.097)	(0.086)	(0.091)
3. Age 45–64	0.900*	0.623*	0.296	0.720*
	(0.358)	(0.226)	(0.176)	(0.229)
4. Age 65+	3.438*	2.408*	1.116	5.219*
	(1.378)	(1.101)	(0.920)	(1.416)
5. Age-adjusted (population weighted average)	**0.692***	**0.587***	**0.330**	**0.936***
	(0.323)	**(0.251)**	**(0.207)**	**(0.289)**
Panel B. Alternative specifications				
1. Separate model for males	0.863	0.711*	0.406	0.968*
	(0.446)	(0.340)	(0.284)	(0.339)
2. Separate model for females	0.461	0.459	0.250	0.887*
	(0.309)	(0.238)	(0.191)	(0.301)
3. With interactions of temperature and precipitation	0.683	0.637	0.386	1.283*
	(0.354)	(0.258)	(0.208)	(0.373)
4. Models for log mortality rate	0.865*	0.663*	0.276	0.879*
	(0.331)	(0.254)	(0.208)	(0.312)
5. Including current and lagged weather variables	0.859	1.152*	0.558	1.255*
	(0.493)	(0.448)	(0.315)	(0.446)
6. Data for 1980 + only	0.103	0.292	0.401	0.852*
	(0.356)	(0.260)	(0.228)	(0.285)
7. Counties with above median annual days >=90° F	1.385*	0.793	0.390	0.892*
	(0.587)	(0.581)	(0.210)	(0.289)
8. Counties with below median annual days <=90° F	0.901*	0.472	0.235	1.020
	(0.360)	(0.297)	(0.291)	(7.138)
9. Missing weather data imputed by multiple imputation	0.547	0.386	0.514*	0.885*
	(0.336)	(0.232)	(0.215)	(0.420)

Notes: The estimates are from fixed-effect regressions by age group weighted by the population count in the relevant age group in a county-year (see equation (5)). With the exception of rows 1–4, the age group-specific estimates are combined into an "age-adjusted" estimate by taking a weighted average of the age-specific estimates, where the weight is the average population in each age group. For each group there are 107,310 county-year observations over the 1968–2002 period. The dependent variable is the annual mortality rate in the relevant age group in a county-year. Temperature exposure is modeled with nine temperature-day bins defined as the number of days in a given temperature category in a county-year. The estimates on the lowest two (coldest) and highest two (hottest) bins are reported. Other control variables include a set of 11 indicator variables capturing the full distribution of annual precipitations. Standard errors are clustered at the county-by-age group level. Starred entries are statistically significant at the 5 percent level. See the text for further details.

Figure 2. We do not report the estimated parameters associated with the precipitation variables, because they are insignficant generally.[14]

There are several important observations to be made from panel A. First, the effect of extreme weather on mortality generally grows with age after infancy. Second,

[14] In addition, the climate change calculations reported below show that the inclusion of the precipitation variables is largely inconsequential.

exposure to extreme temperature significantly increases mortality in all age groups, especially exposure to days with temperature above 90° F.[15] It is also notable that for the largest age group in terms of population (age 1–44), exposure to days with temperature 80–90° F is also significantly associated with excess mortality.

Panel B reports on the robustness of the baseline age-adjusted estimates reported in row 5 of panel A. In rows 1 and 2, the models are estimated separately for men and women, and suggest that the mortality effect of extreme temperature is larger for men than women. The row 3 specification adds interactions between the 9 temperature variables and the 11 precipitation variables and reports the marginal effects evaluated at the sample means. Row 4 considers a model for log mortality rate instead of its level, which means that across the four age groups a total of 19,308 observations are dropped since the log of zero is not defined. The estimates reported here are converted to units of deaths per 100,000 for comparability and are qualitatively similar to the baseline estimates. Row 5 uses separate sets of the nine temperature bins for the current and previous years' daily mean temperatures to allow for the possibility that equation (5) inadequately accounts for the dynamics of the mortality-temperature relationship. The entries report the sum of the current and previous year's temperature coefficients, which indicate modestly larger total impacts.[16]

Rows 6–8 consider specifications that probe the extent to which "adaptation" is able to mitigate the mortality response to extreme temperatures. Specifically, the row 6 estimates are based on a response function that is obtained by fitting equation (5) with post-1980 data. For example in these years, medical technologies are more advanced, air conditioning is more pervasive, and incomes are higher. In row 7, the response function is estimated with data from the counties where the average number of days per year with a mean temperature above 90° F exceeds the national median (0.06 day per year). The idea is that individuals that are more frequently exposed to extremes are likely to have undertaken a series of adaptations to protect themselves against high temperatures in these counties, and these adaptations may resemble the ones that climate change will induce. In this respect, the resulting response functions may better approximate the long-run impacts of climate change on mortality. The entries in both rows reflect the application of the relevant response function to the full sample. The row 6 entries reveal that the impact of cold days is muted in the more recent data, but there is little change in the consequences of exposure to additional hot days. The row 8 entries fail to find much evidence that these counties are better equipped to handle hot days, although the impact of cold days is modestly larger if one ignores issues of precision.

Finally, row 9 considers a different method for imputing missing weather measurements in the station-level data. Rather than only considering stations without

[15] For the 90° F category, the standard error for the all-age estimate that accounts for spatial correlation (bandwidth of 400KM) and serial correlation (up to 3 periods) is 0.416, making the estimated effect marginally significant. We thank Salomon Hsiang (Columbia University) for providing the code.

[16] There is some evidence that individuals acclimate to higher temperatures over time, so consecutive days with high temperatures (i.e., heat waves) may have a different impact on annual mortality than an equal number of hot days that do not occur consecutively. We also estimated models that added a variable for the number of instances of 5 consecutive days of mean daily temperature above 90° F. The resulting estimates were nearly identical to the baseline estimates reported in row 5 of panel A.

missing records, the approach here uses multiple imputation (Donald B. Rubin 1987) to impute the missing daily level measurements for a fixed set of stations between 1968 and 2002.[17,18]

Taken together the evidence in Table 2 strongly supports the notion that exposure to extreme temperatures impacts mortality. Across the various specifications, almost all of the point estimates associated with extreme high temperatures (i.e., more than 90° F) are statistically significant at conventional levels. Further, it is apparent that these alterations to the baseline specification fail to lead to a meaningful change in the estimates of the effect of extreme temperatures on mortality. Some of them modestly increase the point estimate, while others decrease it, but, in light of the standard errors, none of the alternative estimates appear different than the corresponding baseline estimates.[19]

B. *Exploring the Sources of the Relationship between Extreme Temperatures and Mortality Rates*

This subsection explores two potential sources of the relationship between extreme temperatures and mortality rates. First, Table 1 documented that much of the variation in extreme temperatures is concentrated in specific parts of the country. This is especially true for hot temperatures. These areas of the country may have already adapted or undertaken the investments that mitigate the impact of extreme temperatures on mortality, and we explore this possibility. Second, a central feature of our argument is that annual mortality data combined with daily temperature data provides a valid approach to measure the impacts of temperature on mortality even in the presence of an unknown and complicated dynamic relationship between temperature and mortality that is comprised of delayed impacts and forward displacement.

Geographical Heterogeneity.—Table 3 probes the extent of geographical difference in the response functions linking annual mortality and temperature exposure by reporting age-adjusted estimates specific to each of the nine US census divisions. In addition, column 1 reports the average population in each census division between 1968 and 2002, so that its relative importance in the pooled estimates is apparent.

There is substantial heterogeneity in the estimated effects of extreme temperatures on mortality. For example, the annual mortality impact of exposure to one

[17] Following a referee's suggestion, we also estimated models without weighting the regressions by county population. This means that all counties are treated equally irrespective of differences in population. Our view is that this estimation method is inferior to the weighted regressions used in the rest of the paper since population exposure to extreme temperatures will vary across county due to differences in population. Nevertheless, the point estimates from the unweighted specification are: 0.896 (0.352), 0.286 (0.300), 0.064 (0.172), and 0.566 (0.490) corresponding to columns 1a–1d, respectively. It is also notable that the population-weighted models produce smaller standard errors (20 percent smaller on average) and have higher adjusted R-squares (twice as large on average) than the unweighted models.

[18] In order to assess the influence of outliers in the mortality rate distribution, we also estimated the same models on a "trimmed" sample that excluded mortality rates below and above the first and ninety-ninth percentiles. The results are very similar to those in Table 2.

[19] We also estimated separate equations for cause-specific mortality rates. The estimates are reported in online Appendix Table 2. The main result suggests that extreme temperatures significantly increase mortality due to cardiovascular diseases and have no effect on neoplasms.

TABLE 3—ESTIMATES OF THE IMPACT OF EXTREME TEMPERATURES ON AGE-ADJUSTED ANNUAL MORTALITY RATE, BY US CENSUS DIVISION

		Impact on annual mortality rate per 100,000			
	Avg. population (1)	Days < 10° F (2a)	Days 10°–20° F (2b)	Days 80°–90° F (2c)	Days > 90° F (2d)
Census division					
1. New England	12,829,830	−0.969	−0.087	−1.084	23.202
(CT, ME, MA, NH, RI, VT)		(1.206)	(0.756)	(1.193)	(19.650)
2. Middle Atlantic	37,833,310	0.567	1.505*	1.607	−6.251
(NJ, NY, PA)		(0.883)	(0.752)	(1.127)	(9.291)
3. East North Central	42,161,530	1.061*	0.257	−0.032	1.122
(IN, IL, MI, OH, WI)		(0.432)	(0.368)	(0.313)	(3.604)
4. West North Central	17,605,500	1.375*	0.528	0.508	2.056
(IA, KS, MN, MO, NE, ND, SD)		(0.481)	(0.484)	(0.375)	(1.196)
5. South Atlantic	40,143,610	4.800*	3.676*	0.453	5.791*
(DE, DC, FL, GA, MD, NC, SC, VA, WV)		(1.990)	(1.558)	(0.503)	(2.269)
6. East South Central	14,979,270	0.197	0.575	0.016	3.938*
(AL, KY, MS, TN)		(2.252)	(0.937)	(0.346)	(1.482)
7. West South Central	25,526,750	−3.077	−0.852	−0.205	−0.014
(AR, LA, OK, TX)		(2.020)	(0.959)	(0.268)	(0.349)
8. Mountain	12,924,740	0.866	0.400	0.447	1.234*
(AZ, CO, ID, NM, MT, UT, NV, WY)		(0.549)	(0.463)	(0.380)	(0.530)
9. Pacific	34,095,370	−0.270	0.567	0.643	0.688
(CA, OR, WA)		(1.423)	(0.682)	(0.429)	(0.745)
10. Population weighted average of division-specific estimates	**—**	**−0.792**	**1.004**	**0.399**	**1.932**
		(1.234)	**(0.799)**	**(0.534)**	**(3.868)**
11. Baseline estimate (from Table 2)	**—**	**−0.692***	**0.587***	**0.330**	**0.936***
		(0.323)	**(0.251)**	**(0.207)**	**(0.289)**

Notes: The estimates are from fixed-effect regressions estimated separately by age group and US census division. The age group-specific estimates are then combined into an "age-adjusted" estimate by taking a weighted average of the age-specific estimates, where the weight is each age group's share of the total population. The dependent variable is the annual mortality rate in the relevant age group in a county-year. Temperature exposure is modeled with nine temperature-day bins defined as the number of days in a given temperature category in a county-year. The estimates on the lowest two (coldest) and highest two (hottest) bins are reported. Standard errors are clustered at the county-by-age group level. Starred entries are statistically significant at the 5 percent level. See the text for further details.

day where the mean temperature exceeds 90° F ranges from −6.3 to 23.2 additional deaths per 100,000, although these two point estimates have large standard errors. It is also notable that the estimated mortality impact of extreme high temperature is positive in seven out of nine divisions and statistically significant in three out of nine divisions (South Atlantic, East South Central, and Mountain). Likewise, the estimated mortality impact of extreme low temperature is statistically significant in three out of nine divisions (East North Central, West North Central, and South Atlantic).[20]

[20] The last two rows of Table 3 report aggregate (i.e., national) estimates. Row 10 aggregates the division-specific estimate by taking a weighted average where the weights correspond to the share of total population in each division. Row 11 reproduces the baseline estimates from Table 2. It is evident that the national estimates obtained from the division-specific ones are larger in magnitude than the ones from the pooled model that restricts the effect of temperature (and of all the other model variables) to be the same across all counties. However their standard errors are also larger, so the results in row 10 should be interpreted cautiously.

Finally, the null that the parameters associated with the 4 temperature variables reported in this table are equivalent for the 4 age groups across the 9 divisions is rejected at conventional levels 9 out of 16 times. It is noteworthy that this hypothesis is rejected for all four age groups for the $> 90°$ F temperature variable.[21]

A frequent explanation for the differences in response functions across divisions is that hotter places (e.g., Houston) are better adapted to respond to extremely hot temperatures, likewise colder places (e.g., Minneapolis) are more likely to have made the investments necessary to protect themselves against extremely cold temperatures (see, e.g., Basu and Samet 2002). We investigated the former by regressing the division-specific parameter estimate on the variable for the number of days with a mean temperature exceeding 90° F against the division-specific mean of this variable. The revealed relationship between these two variables is weak and statistically insignificant, regardless of whether the regression is weighted by the inverse of the standard error of the parameter estimate. We also found evidence of a weak relationship between the division-specific parameter estimates on the variable for the number of days with a mean temperature below 10° F against the division-specific mean of this variable.

Overall, these two tests, each based on just nine data points, fail to produce empirical support for the view that differential adaptation to hot and cold weather explains the heterogeneity across divisions. So although there appears to be variation in the response functions across the divisions, these data fail to find that it is related to climate or baseline temperatures. This may simply reflect that cooling and heating systems were already pervasive throughout the United States in this period.

Annual Data versus Daily Data.—A striking result in Table 2 is the small magnitude of the estimated impact of extreme temperature on mortality, relative to conventional wisdom about high temperatures. Our baseline estimates imply that one additional day with mean temperature above 90° F, leads to about 1 extra *annual* death per 100,000 population. These are small in comparison with the estimates derived from recent heat waves. For example, the Cook County coroner estimated that there were 465 heat-related deaths in the 1995 Chicago heat wave during the week of July 14–20, and S. Whitman et al. (1997) concluded that the heat wave lead to 739 "excess deaths." If correct, these impacts are enormous in a location with a population of about 5 million. Indeed, the conventional wisdom on the effect of exposure to high temperatures on mortality is largely based on event studies of heat waves such as the ones described above.

Unfortunately, these studies and their results are difficult to compare to ours for several reasons. First, they are typically based on data from a single city or just a handful of cities. Second, they are usually based on a limited number of years. Third, they typically do not model the effect of temperature on mortality explicitly. Rather, they frequently compare observed mortality in a location during a heat

[21] We tested the equality of the temperature effects across the nine US census divisions for each temperature-day category and each age group. For temperature-days below 10° F (column 2a), the *p*-values from these tests are 0.316, 0.255, 0.003, and 0.002 (corresponding, respectively, to infants, age 1–44, age 45–64, and age 65+). For the remaining 3 columns of Table 3, the *p*-values are 0.044, 0.002, 0.030, 0.210 (column 2b); 0.115, 0.308, 0.193, 0.055 (column 2c); and 0.001, 0.001, 0.001, 0.002 (column 2d).

wave or a period of extreme temperature with the baseline mortality, defined as the observed mortality in the same period and location in the several years prior to the event. Thus, the idea of a response function like the one that is estimated here is largely foreign in this literature.

The study that is closest to ours is Deschênes and Moretti (2009). Their study finds that there is a complicated dynamic relationship between daily temperatures and mortality rates. For example, they find that the number of deaths immediately caused by a period of very high temperatures is at least partially compensated for by a reduction in the number of deaths in the period immediately subsequent to the hot day or days. This pattern is called forward displacement or "harvesting," and it appears to occur because heat affects individuals who were already very sick and would have died in the near term. Additionally in the case of cold days, the effect is the opposite, as the total effect of a cold day grows over time and this is referred to as "delayed impacts."

Our study uses annual mortality data, rather than the daily mortality data that Deschênes and Moretti (2009) utilize, for several interrelated reasons. First, Deschênes and Moretti (2009) show that the full impact of a given day's temperature will take at least 30 days to manifest itself fully, so it is necessary to include many lags of the temperature variables in equations for daily mortality. Second, the temperature-mortality relationship is highly nonlinear and requires the inclusion of several temperature variables. Due to Deschênes and Moretti's (2009) result, it would be necessary to include at least 30 lags for each of these variables in daily mortality equations. Third, the MCOD county-level daily mortality is only available for the entire United States from 1972–1988, compared to 1968–2002 for the annual CMF data, which poses a challenge for the estimation of some of the very high and very low temperature categories (recall Table 1).[22] Since most deaths in a calendar year will be functions of the temperature realization in the same year, the combination of annual mortality data and aggregated daily temperature data should be sufficient to flexibly capture the full dynamic relationship between temperature and mortality (recall the results in Table 2, panel B, row 5 that include the current and lagged year's temperature variables are qualitatively similar to the baseline results).

It should be noted that the estimates of the response function in this paper are not directly comparable with those in Deschênes and Moretti (2009). For one, the main objectives of both papers differ. Deschênes and Moretti (2009) is mostly concerned with the determination and estimation of the appropriate adjustment window in which temperature shocks affects daily mortality. The specification of the models also greatly differs in Deschênes and Moretti (2009) and in this paper. The models in Deschênes and Moretti (2009) rely on a single temperature category (e.g., days colder than 30°F and warmer than 80°F) as opposed to models with several temperature-day categories as in this paper. Finally, the main specification in Deschênes and Moretti (2009) included a full set of county-by-year-by-month effects from

[22] Furthermore, the MCOD data with the date of death appear much noisier. For example, the estimation of equation (5) with the MCOD for the years 1972–1988 produces standard errors that are about 1.8 times larger than the standard errors on the same variables from the estimation of the same equation in the same years with the CMF data.

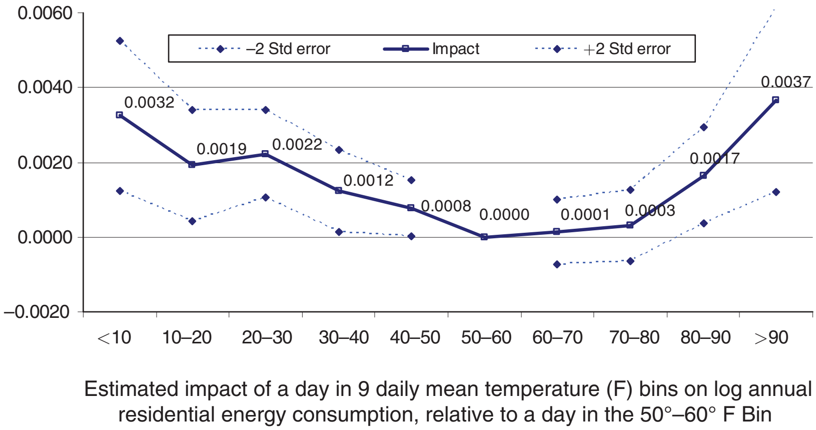

Figure 3. Estimated Relationship Between Log Residential Energy Consumption and Average Daily Temperature

Notes: Figure 3 plots the estimated response function between log annual residential energy consumption and daily mean temperatures. This is obtained by fitting equation (6) on our sample of 1,715 state-year observations. The response function is normalized with the 50º–60ºF category set equal to 0 so each θ_j corresponds to the estimated impact of an additional day in bin j on residential QBTU relative to the log residential energy QBTU associated with a day where the temperature is between 50º–60ºF. The figure also plots the θ_js plus and minus one standard error of the estimates. The numbers above the response function correspond to the point estimates associated with each temperature bin.

which it is difficult to draw comparisons to the models used in this paper, which are based on annual data. Nevertheless, the 95 percent confidence intervals of the estimates of the effect of high temperature days (i.e., 80–90°F and >90°F) from daily mortality models (with 30 lags) overlaps with the 95 percent confidence intervals of the estimates from annual mortality data derived from the 1972–1988 MCOD.

C. *Estimates of the Impact of Temperature on Residential Energy Consumption*

We now turn to an analysis of the effect of inter-annual variation in temperature on log residential energy consumption. Specifically, this subsection reports on the fitting of versions of equation (6) to the state-by-year data on residential energy consumption from the EIA. Figure 3 plots the estimated response function linking log residential energy consumption, and the ten daily temperature bin variables, after adjustment for the equation (6) covariates. The coefficients report the estimated impact of an additional day in bin j on annual energy consumption, relative to energy consumption on a day where the temperature is in the 50°–60°F range. The figure also plots the coefficients plus and minus two standard errors, so their precision is evident.

Overall, the response function has a U-shape, so that the hottest and coldest days are the highest energy consumption ones. Temperature-days in the highest two categories (80–90°F and >90°F) and the lowest four categories (30–40°F and the three categories below) are associated with statistically significant increases in residential energy consumption. In terms of magnitude, the estimated impacts are similar

Table 4—Estimates of the Impact of Extreme Temperatures on Aggregate Annual Residential Energy Consumption

	Average aggregate annual residential energy consumption (QBTU)	Impact on log annual residential energy consumption (QBTU)					
		Days < 10° F (1a)	Days 10°–20° F (1b)	Days 80°–90° F (1c)	Days >90° F (1d)	HDD 65 (hnd.) (2a)	CDD 65 (hnd.) (2b)
Panel A. Pooled estimates:	16.62	0.0032*	0.0019*	0.0017*	0.0037*	0.0053*	0.0105*
		(0.0010)	(0.0007)	(0.0006)	(0.0012)	(0.0017)	(0.0025)
Panel B. Census-division specific estimates:							
1. New England	0.95	0.0024	0.0039	0.0014	0.1012	0.0047	0.0014
(CT, ME, MA, NH, RI, VT)		(0.0027)	(0.0021)	(0.0040)	(0.1323)	(0.0044)	(0.0163)
2. Middle Atlantic	2.48	0.0096*	0.0084*	−0.0044	0.0285	0.0228*	0.0057
(NJ, NY, PA)		(0.0046)	(0.0033)	(0.0023)	(0.0257)	(0.0071)	(0.0080)
3. East North Central	3.38	0.0014	−0.0004	0.0003	−0.0016	0.0024	0.0075
(IN, IL, MI, OH, WI)		(0.0015)	(0.0011)	(0.0014)	(0.0109)	(0.0020)	(0.0045)
4. West North Central	1.38	0.0030*	0.0017	0.0005	−0.0022	0.0050*	0.0040
(IA, KS, MN, MO, NE, ND, SD)		(0.0010)	(0.0011)	(0.0011)	(0.0022)	(0.0015)	(0.0036)
5. South Atlantic (DE, DC,	2.80	0.0108*	0.0027	0.0000	0.0070	0.0116*	0.0066*
FL, GA, MD, NC, SC, VA, WV)		(0.0053)	(0.0019)	(0.0007)	(0.0061)	(0.0017)	(0.0022)
6. East South Central	1.14	0.0130*	0.0073*	0.0024	0.0207*	0.0128*	0.0145*
(AL, KY, MS, TN)		(0.0050)	(0.0025)	(0.0015)	(0.0093)	(0.0049)	(0.0055)
7. West South Central	1.78	0.0166*	0.0014	0.0018	0.0010	0.0069*	0.0139*
(AR, LA, OK, TX)		(0.0068)	(0.0041)	(0.0011)	(0.0018)	(0.0034)	(0.0044)
8. Mountain (AZ, CO,	0.84	0.0056*	0.0022	0.0003	0.0026*	0.0073*	0.0066*
ID, NM, MT, UT, NV, WY)		(0.0018)	(0.0023)	(0.0010)	(0.0013)	(0.0015)	(0.0033)
9. Pacific	1.87	−0.0103	0.0185*	−0.0029	0.0343*	0.0046*	−0.0056
(CA, OR, WA)		(0.0130)	(0.0060)	(0.0023)	(0.0163)	(0.0023)	(0.0051)
10. p-value from F-tests of equality-Across US census divisions	—	0.001	0.001	0.001	0.088	0.003	0.001

Notes: The estimates are from fixed-effect regressions based on a sample of 1,715 state-year observations (see equation 6). Each model includes state fixed-effects and census division-by-year fixed-effects. The dependent variable is the log of the total residential energy consumption in a state-year (in QBTU). Control variables include quadratics in population, state GDP, and their interactions, as well as a set of 11 indicator variables capturing the full distribution of annual precipitation. In columns 1a–1d, temperature exposure is modeled with nine temperature-day "bins" defined as the number of days in a given temperature category in a state-year. The estimates on the lowest two (coldest) and highest two (hottest) bins are reported. The entries in columns 2a and 2b are from a separate regression where temperature is modeled in hundreds of heating degree days (HDD) and cooling degree days (CDD), both of which are calculated with a base of 65° F. Standard errors are clustered at the state level. Starred entries are statistically significant at the 5 percent level. See the text for further details.

at the upper and lower tail of the daily temperature distribution; temperature-days below 10°F and above 90°F are associated with 0.3 percent–0.4 percent increases in annual residential energy consumption. The results are summarized in Table 4, and its layout is similar to Tables 2 and 3. Panel A reports pooled estimates, while panel B reports estimates specific for each US census division. The first column reports the average aggregate annual residential energy consumption, which is 16.6 quadrillion BTUs, QBTU over the 1968–2002 period.

The remaining columns report the results from two separate versions of equation (6), where the dependent variable is log of the total residential energy consumption measured at the state by year level. Again, only the point estimates corresponding to the lowest two (coldest) and highest two (hottest) exposure bins are reported in columns 1a–1d, although all models include the nine temperature bins described earlier.

In addition, we also report estimates based on cooling and heating degree days (both use a base of 65°F) since it is the dominant approach in the literature (e.g., Robert F. Engle et al. 1986). For ease of interpretation, the point estimates reported in columns 2a and 2b are scaled in hundreds of cooling and heating degree days. An additional hundred cooling degree-days (roughly one-tenth of the annual national average) increases energy consumption by 1.1 percent or 0.182 QBTU, while an additional hundred heating degree-days (roughly one-fiftieth of the annual national average) increases consumption by 0.5 percent or 0.087 QBTU.

Panel B reveals tremendous heterogeneity in the response function for residential energy consumption, although statistical precision is an issue.[23] This is observed in both the degree-days and the "bins" specification. Focusing on the estimates in columns 1a–1d, the impact of days in the highest temperature bin on residential energy consumption is statistically significant in only three of nine divisions (East South Central, Mountain, and Pacific). The three divisions where the response to high temperature days is the largest are New England, Middle Atlantic, and Pacific, with impacts ranging from 3 to 10 percent, although none of these is statistically significant at the 5 percent level. The estimated effect of cooling degree-days is statistically significant at the 5 percent level for the South Atlantic, East South Central, West South Central, and Mountain divisions.

The differential impact of days in the lowest temperature category on residential energy consumption across division is better determined in the data, reflecting the greater frequency of these days. The point estimates associated with the lowest temperature category (column 1a) are statistically significant in 6 of 9 divisions (Middle Atlantic, South Atlantic, West North Central, East South Central, West South Central, and Mountain) and correspond to impacts of 0.3 percent to 1.7 percent. In the heating degree-days specification, the impacts are significant in seven out of nine divisions.

An important question about the generalizability of these results to the future is the degree to which there are currently differences in adaptation across the United States that might change in responses to climate change. For example, higher temperatures could increase the rate of penetration of air conditioning in areas that currently have few extremely hot days. We assessed this possibility in two ways. The cruder approach is a test of whether the division-specific estimates are equal. The bottom row of Table 4 reveals that this null is rejected at conventional levels for the parameters associated with the temperature variables, except for the $>90°$ F variable that has a p-value of 0.088.

The second approach is based on the idea that hot days should cause greater increases in energy consumption in parts of the country where high temperatures are relatively frequent. We tested this idea by regressing the division-specific $>90°$ F parameters on the division-specific mean number of days in that category annually. This nine data point regression fails to find evidence that the frequency of hot days is related to the energy consumption responsiveness of hot days. Indeed, the parameter estimate has a counter-intuitive negative sign and is very poorly determined (t-statistic less than 1), regardless of whether it is weighted by the inverse of the parameter estimate's standard error. This finding is consistent with the possibility

[23] Anin Aroonruengsawat and Maximilian Auffhammer (2011) also report important geographical differences in the effect of extreme temperatures on residential consumption using zip code level data for the state of California.

that air conditioning or other cooling technologies are already pervasive in the United States and/or that households in warmer areas have undertaken other investments that mitigate the need for greater energy consumption on hot days.[24]

We believe that this analysis of the relationship between temperature and energy consumption has made at least two advances on previous research. First, it reveals that the standard approach of modeling energy consumption with heating and cooling degree-days obscures the nonlinear increase in energy consumption at extremely high temperatures. Second, these estimates are relevant to the entire country. In contrast, much of the previous literature is based on data from a specific city or state, which leaves questions of external validity unanswered (e.g., Guido Franco and Alan H. Sanstad 2006).

D. *Predicting the Impact of Climate Change on Mortality Rates and Residential Energy Consumption*

The relationships between temperature and mortality and energy consumption that this paper has documented can be combined with predictions about climate change to develop measures of some of the costs of climate change. As Figure 1 demonstrated, the state-of-the-art climate models predict dramatic increases in the number of extremely hot days, especially in the 80–90° F and >90° F bins, when mortality and energy consumption are relatively high. Importantly, these increases are largely predicted to be offset by decreases in the number of days in the middle of the temperature distribution, where mortality rates and energy consumption are the lowest. Thus, it seems that climate change will cause the United States to exchange relatively low cost days for ones with higher social costs.

We now turn to a more precise calculation of the predicted mortality impacts of climate change. Table 5 summarizes the results from the estimation of separate versions of equation (5) for the four age groups using the error-corrected Hadley 3 A1FI predictions. Here and throughout the rest of the analysis (except Table 7), the climate change predictions are for the end of the century and this is defined as the average of the predictions for the years 2070–2099. The final row sums the estimated impacts to provide all age estimates. The empirical estimates underlying the calculations in Table 5 are based on the regression results reported in panel A of Table 2.

The entries are based on calculations of the estimated impact of a climate change scenario on annual US mortality for each age group. For example, the estimated impact of predicted temperature changes on a given county and age group is calculated as follows:

$$M_{ca} = MeanPopulation_{ca} \times \sum_{j} \left(\hat{\theta}_{aj}^{TMEAN} \bullet \Delta TMEAN_{cj}\right). \tag{7}$$

[24] We conducted the analogous exercise with the division-specific parameters on the variable for the number of days where the temperature is less than 10° F. This regression also produced a "perversely" signed relationship that has a t-statistic of roughly 1.55 when the regression is weighted by the inverse of the parameter estimate's standard error.

Table 5—Estimates of the Impact of Climate Change on Annual Mortality, By Age Group and by Segment of the Temperature Distribution, Based on Error-Corrected Hadley 3 A1FI Model

	Impact of change in days with temperature			Total temperature impact	Temperature and precipitation impact	% Change in annual mortality	Annual change in life-years
Age group:	<20° F (1a)	20°–80° F (1b)	> 80° F (1c)	(2)	(3)	(4)	(5)
Infants	156	−1,097	4,293*	3,352	3,318	7.8	−245,134
	(254)	(626)	(1,992)	(1,761)	(1,767)	(4.1)	
1–44	3,784*	−14,055	*26,745	*8,906	8,461	4.7	−453,510
	(1,445)	(6,635)	(9,997)	(7,060)	(7,099)	(3.9)	
45–64	−3,362*	−9,114*	19,738*	7,262	6,985	1.7	−172,180
	(1,112)	(2,908)	(6,455)	(5,150)	(5,197)	(1.3)	
65+	−7,704*	−20,440*	71,986*	43,843*	44,266*	3.0*	−418,183
	(3,041)	(6,937)	(22,286)	(17,429)	(17,504)	(1.2)	
All ages	**−14,693***	**−44,706***	**122,763***	**63,363***	**63,030***	**3.0***	**−1,289,006**
	(5,851)	**(17,107)**	**(40,731)**	**(31,400)**	**(31,567)**	**(1.5)**	

Notes: The estimates are from fixed-effect regressions by age group, for ten temperature-day bins and the other control variables as reported in rows 1–5 in Table 2. Standard errors (in parentheses) are clustered at the county-by-age group level and take climate change predictions as constants. Starred entries are statistically significant at the 5 percent level. Years of life lost based on 1980 life table estimates. See the text for further details.

That is, the predicted change in the number of days in each temperature cell $(\Delta TMEAN_{cj})$ is multiplied by the corresponding age group-specific impact on mortality $\left(\hat{\theta}_{aj}^{TMEAN}\right)$, and then these products are summed. This sum is then multiplied by the average population for that age group in that county $(MeanPopulation_{ca})$ over the sample period. The impacts for a given age group are then summed over all counties. This sum is the national age group-specific estimate of the change in annual mortality. It is straightforward to calculate the standard error (reported below the estimate in parentheses), since the estimated mortality change is a linear function of the parameters.

Columns 1a–1c report this calculation for the bottom two temperature categories (i.e., $< 10°$F and 10°–20°F), the middle five temperature categories, and top two temperature categories (i.e., 80°–90°F and >90°F). The entries are the change in annual fatalities, based on the assumption that the size and geographical distribution of the population remain constant. Column 2 reports the total temperature impact and column 3 adds in the impact of the relatively inconsequential change in precipitation (which is calculated analogously to the method outlined in equation (7)). Column 4 reports the estimated percentage change in the annual mortality rate (and its standard error), which is calculated as the ratio of the change in the age group's annual mortality rate due to predicted climate change to its overall annual mortality rate.

Column 5 reports the change in life-years due to predicted climate change for each age category (again assuming a constant population and geographical distribution). This entry is the product of the predicted change in annual fatalities and the residual life estimate for each age group, which is evaluated in the middle of the age range and taken from the 1980 Vital Statistics;[25] the negative values correspond to

[25] The residual life estimates for the four age categories are 73.9 (infants), 53.6 (1–44), 24.7 (45–64), and 9.4 (>65).

losses of life-years. This calculation may overstate the change in life-years, because affected individuals are likely to have shorter life expectancies than the average person. Nevertheless, these entries provide a way to capture that fatalities at young ages may have greater losses of life expectancy than those at older ages.

The error-corrected Hadley 3 A1FI results suggest that climate change would lead to approximately 63,000 additional deaths annually in the United States (as reported in the bottom row of column 3), which is equivalent to a 3 percent increase in the annual mortality rate. It is apparent that the overall effect is driven by statistically significant increases in mortality due to the predicted increased incidence of very hot days (122,763 extra deaths in column 1c) and predicted decreases in mortality due to fewer moderate days (44,706 fewer deaths in column 1b) and very cold days (14,693 fewer deaths in column 1a).[26] As a result, the overall effect of 63,030 extra annual deaths is comprised of statistically significant increases and decreases in annual mortality and is itself marginally statistically significant at the conventional significance levels. The life-years calculation implies an annual loss of roughly 1.3 million life-years.

Two other findings bear highlighting. First, the predicted increase in the annual mortality rate is greatest for infants; the estimates suggest a 7.8 percent increase in the infant mortality rate or 3,318 additional infant deaths per year. This result is not surprising, because infants' thermoregulatory systems are not yet fully functional (Robin Knobel and Diane Holditch-Davis 2007). Second, the percent impacts on the other age categories are smaller and only the increase among the 65 and older category would be judged statistically significant at conventional levels.

Table 6 reports predicted impacts of climate change on annual mortality and aggregate residential energy consumption by US census division. For reasons of precision, we use the national estimates of the mortality and residential energy responses to temperature, rather than the division-specific ones. However, the predicted changes in the temperature and precipitation distributions are specific to each division. The estimates are reported for the changes occurring in the lowest and highest temperature bins (columns 1 and 2, and 4 and 5, respectively) as well as the overall impact that accounts for the full changes in the temperature and precipitation distributions (columns 3 and 6).

The first row shows the pooled estimates, which in the case of annual mortality have already appeared in Table 5. In terms of aggregate residential energy consumption, the pooled estimates indicate that climate change will lead to a net increase of 1.9 QBTU, with a standard error of 1.2.

The remaining rows of Table 6 report the predicted impacts of climate change on annual mortality and aggregate residential energy consumption by US census division. The overall mortality impacts are reported in column 3. For the first four divisions (New England—West North Central), which are all relatively cold-weather locations, the predicted net effect of climate change is a relatively small and statistically insignificant increase in annual mortality. The predicted change in mortality is positive for the remaining five divisions and is statistically significant in four of

[26] As shown in Figure 2, all the point estimates associated with temperatures below 50° F are statistically significant.

Table 6—Estimates of the Impact of Climate Change on Annual Mortality and Residential Energy Consumption, By Census Division and Segment of Temperature Distribution

	Impact on annual mortality			Impact on annual residential energy consumption		
	Days < 20° F (1)	Days > 80° F (2)	Overall impact (3)	Days < 20° F (4)	Days > 80° F (5)	Overall impact (6)
Pooled estimates:	**−14,693***	**122,763***	**63,030***	**−0.537***	**3.392***	**1.877**
	(5,851)	**(40,731)**	**(31,567)**	**(0.162)**	**(1.058)**	**(1.180)**
Estimates by census division:						
1. New England	−1,601*	4,604*	538	−0.120*	0.277*	0.035
(CT, ME, MA, NH, RI, VT)	(636)	(1,731)	(1,543)	(0.037)	(0.087)	(0.136)
2. Middle Atlantic	−3,032*	18,734*	6,201	−0.026*	0.208*	0.092
(NJ, NY, PA)	(1,211)	(6,432)	(5,089)	(0.009)	(0.065)	(0.078)
3. East North Central	−5,383*	18,583*	5,084	−0.088*	0.316*	0.116
(IN, IL, MI, OH, WI)	(2,148)	(6,326)	(5,221)	(0.027)	(0.098)	(0.120)
4. West North Central (IA, KS,	−2,927*	10,021*	3,759	−0.179*	0.499*	0.202
MN, MO, NE, ND, SD)	(1,185)	(3,167)	(2,692)	(0.053)	(0.157)	(0.183)
5. South Atlantic (DE, DC, FL,	−483*	23,508*	14,673*	-0.022*	0.751*	0.557*
GA, MD, NC, SC, VA, WV)	(196)	(8,254)	(5,855)	(0.008)	(0.234)	(0.230)
6. East South Central	−331*	9,546*	5,721*	−0.010*	0.361*	0.274*
(AL, KY, MS, TN)	(134)	(3,057)	(2,372)	(0.003)	(0.114)	(0.107)
7. West South Central	−180*	20,550*	15,514*	−0.006*	0.437*	0.358*
(AR, LA, OK, TX)	(74)	(6,475)	(5,909)	(0.002)	(0.146)	(0.137)
8. Mountain (AZ, CO, ID,	−642*	4,528*	1,948	−0.081*	0.405*	0.154
NM, MT, UT, NV, WY)	(259)	(1,505)	(1,222)	(0.025)	(0.126)	(0.161)
9. Pacific	−115*	12,687*	9,592*	−0.003*	0.138*	0.089
(CA, OR, WA)	(46)	(4,837)	(4,198)	(0.001)	(0.043)	(0.059)

Notes: Columns 1–3 report the predicted age-adjusted change in annual mortality due to climate change, based on the results from fitting equation (5) in panel A of Table 2, and include the effect of predicted changes in temperature and precipitation distributions (precipitation effects in column 3 only). Columns 4–6 present the predicted change in annual residential energy consumption due to climate change, based on the results from fitting equation (6) which is reported in row A of Table 5, and include the effect of predicted changes in temperature and precipitation distributions (precipitation effects in column 6 only). Standard errors are clustered at the county-by-age group level (annual mortality estimates) or state level (annual residential energy consumption estimates) and take the error-corrected Hadley 3 A1FI climate change predictions as constants. Starred entries are statistically significant at the 5 percent level. See the text for further details.

them (South Atlantic, East South Central, West South Central, and Pacific). Indeed, these four divisions account for nearly 70 percent of the total increase in mortality.

The results for residential energy consumption by division generally exhibit the same patterns as those for annual mortality. The overall impacts (column 6) suggest that residential energy demands will increase in most divisions due to climate change. Just as with mortality, some of the largest increases are projected to occur in the South Atlantic, East South Central, and West South Central divisions. This finding of larger impacts in the South is due to the regional differences in projected changes in climate (since the national response function is used for all regions).

VI. Interpretation

Optimal decisions about climate change policies require estimates of individuals' willingness to pay to avoid climate change over the long run. Previous research has suggested that human health is likely to be a big part of these costs. This section assesses the magnitude of the estimated impacts in the United States and discusses some caveats to this exercise.

The central tendency of the baseline mortality estimates is that by the end of the twenty-first century, the overall mortality rate will increase by about 3 percent with the error-corrected Hadley 3 A1FI predictions. To put this estimate in some context, the US age adjusted death rate for both genders has decreased from 1,304.5 to 832.7 per 100,000 between 1968 and 2002, which is a decline of approximately 1 percent per year. Thus, even if the point estimates are taken literally, the climate change induced increase in mortality is roughly equivalent to losing just three years of typical improvement in longevity over this century.

As an alternative, Table 7 presents the present discounted value of the expected welfare loss associated with predicted changes in climate during the twenty-first century. Columns 1a and 1b present the predicted changes in life-years and energy consumption (measured in quads) over the periods 2010–2039, 2040–2069, 2070–2099, and 2010–2099. The standard errors are reported in parentheses below the estimated impacts.

The total loss in life-years is about 68 million life years between 2010 and the end of the century. The increase in energy consumption is roughly 90 QBTU. The changes in life-years and energy consumption are largest in the 2070–2099 period, because the predicted temperature changes are largest during these years.

Columns 2a and 2b monetize these changes with two different sets of assumptions. In column 2a, we assume the value of a life year is $100,000 and the cost of a quad of energy consumption is $10 billion (2010$).[27] The calculations in column 2b are based on the assumptions that real per capita income grows by 2 percent per year and the elasticity of the value of a life-year with respect to income is 1.6, which is consistent with Dora L. Costa and Matthew E. Kahn (2004). Further, we assume that real energy prices increase by 5 percent annually. Thus, the current valuations of $100,000 per life-year and $10 billion per quad increase such that the respective valuations are approximately $1,032,000 per life-year and $863.3 billion per quad in 2099. In both columns, we assume a discount rate of 3 percent.

Column 2a indicates that the present discounted values of these climate change induced losses over the remainder of the twenty-first century are about $1.5 trillion. It is also noteworthy that this estimate has an associated *t*-statistic of roughly 1.5, so in this case the null of zero damages is not rejected at conventional significance levels. The losses with the non-error-corrected Hadley 3 A1FI and CCSM 3 A2 predictions are larger. They are $4.0 trillion and $2.1 trillion, respectively, and the null hypothesis of a zero effect is rejected in both cases.[28]

The entries in column 2b are necessarily larger, which highlights that these calculations are sensitive to the assumptions about the evolution of the value of a statistical life-year and energy prices (as well as the discount rate). A judgment about the magnitude of these larger estimates requires decisions on the present value of future

[27] This valuation of a life-year is roughly consistent with Ashenfelter and Greenstone's (2004) estimate of the value of a statistical life. The average cost of a quad in 2010 dollars between 1990 and 2000 is $10 billion.

[28] All standard errors in Table 7 treat the climate change predictions as nonstochastic.

Table 7—Estimates of the Present Discounted Value of the Health Related Welfare Costs of Climate Change by 30 Year Periods

	Change in		Present discounted value of welfare loss ($2010 billion)	
	Life-years (million) (1a)	Energy consumption (QBTU) (1b)	Assume constant energy prices and value of a statistical life year (2b)	Assume annual increases in energy prices and value of a statistical life year (2b)
Hadley 3 A1FI, error-corrected				
2010–2039	−5.95	5.43	$403.6	$664.1
	(3.39)	(5.38)	(242.9)	(409.6)
2040–2069	−21.37	28.63	$631.0	$2,555.9
	(12.70)	(18.88)	(374.9)	(1,564.3)
2070–2099	−40.56	56.90	$495.4	$5,345.7
	(23.98)	(35.32)	(298.5)	(3,218.6)
Total	−67.88	90.96	$1,530.0	$8,565.7
	(39.71)	(59.25)	(903.2)	(5,155.5)
Hadley 3 A1FI, uncorrected				
Total	−122.13*	184.64*	$4,037.3*	$15,659.2*
	(52.68)	(50.98)	(1,455.8)	(6,044.8)
CCSM 3 A2				
Total	−63.07	81.39	$2,133.1*	$7,670.1
	(31.69)	(41.00)	(948.6)	(3,901.9)

Notes: Column 1a reports the predicted change in life-years over different periods. The total rows refer to the years 2010–2099. Column 1b reports the predicted change in energy consumption in quads. The column 2a calculations assume that the value of a statistical life year is $100,000 and that the cost of a quad of energy is $10 billion ($2010). The column 2b calculations assume that energy prices increase by 5 percent real annually. Additionally, they assume that real per capita income grows by 2 percent per year and following Costa and Kahn (2004) that the elasticity of the value of a statistical life year with respect to income is 1.6. All present value calculations use a 3 percent discount rate. Standard errors are reported in parentheses. Starred entries are statistically significant at the 5 percent level. See the text for further details.

income growth. For this reason, we emphasize the column 2a estimates. In addition, the estimates in column 2a generally have tighter confidence intervals.

There are a number of caveats to these calculations, and to the analysis more generally, that bear noting. First, the effort to project outcomes at the end of the century requires a number of strong assumptions, including that the climate change predictions are correct, relative prices (e.g., for energy and medical services) will remain constant, the same energy and medical technologies will prevail, and the demographics of the US population (e.g., age structure) and their geographic distribution will remain unchanged. These assumptions are strong, but their benefit is that they allow for a transparent analysis based on the available data.

Second, there is considerable uncertainty about climate predictions, and this paper's estimates are unable to account for this uncertainty. The state of climate modeling has advanced dramatically over the last several years, but there is still much to learn, especially about the role of greenhouse gases on climate (Thomas R. Karl and Kevin E. Trenberth 2003). Thus, the Hadley 3 A1FI and CCSM 3 A2 predictions should be conceived of as two realizations from a superpopulation of models and scenarios. The sources and magnitude of uncertainty in these models and scenarios are unclear, so they cannot readily be incorporated into the estimates of the impacts of climate change. Nevertheless, the use of two sets of daily business as usual climate change predictions provides some sense of the variation.[29]

Third, the life-years calculation assumes that the individuals whose lives are affected by the temperature changes had a life expectancy of 78.6 for women and 71.2 for men. It seems plausible and perhaps even likely that the individuals that die due to higher temperatures had below average life expectancies, in which case the estimated loss of life-years is too high. Fourth, the predicted impacts on mortality totals (i.e., not rates) implicitly assume that the total population remains constant. Those could be adjusted when combined with predictions about future population levels.

Fifth, these calculations are unlikely to capture the full impact of climate change on health. In particular, there may be increases in the incidence of morbidities due to the temperature increases. Additionally, there are a series of indirect channels through which climate change could affect human health, including greater incidence of vector borne infectious diseases (e.g., malaria and dengue fever). Further, it is possible that the incidence of extreme events would increase, and these could affect human health (Kerry Emanuel 2005). However, this study is not equipped to shed light on these issues.

Sixth, and perhaps most importantly, the theoretical section highlighted that our estimates likely overstate the increase in mortality and energy consumption due to climate change. Our identification strategy relies on inter-annual fluctuations in weather, rather than a permanent change, and there are a number of adaptations that cannot be undertaken in response to a single year's weather realization. For example, permanent climate change is likely to lead to institutional adaptations (e.g., improvements in hospitals' ability to treat heat related illnesses) and perhaps even migration. Our approach fails to account for these adaptations. Although some of these adaptations may be costly, individuals will only undertake them if they are less costly than the alternative. For this reason, our approach is likely to overstate the part of the health costs of climate change that we can estimate.

VII. Conclusions

There are several broader implications of this research. First, the age group-specific mortality and residential energy consumption response functions are not specific to any climate model. They reflect the current level of infrastructure, physiology, medical technology, adaptation technologies, and possibly many other factors. As such, they can be used today to understand the complex relationship between temperature deviations and human well-being that prevails today.

Second, this paper has demonstrated that it is possible to develop harvesting and delayed-impact resistant estimates of the impacts of weather on mortality by combining *annual* mortality data with *daily* weather data. In principle, this approach can be applied to other settings where there is an unknown dynamic relationship

[29] We are unaware of a single previous impact study that utilizes *daily* predictions from even a single model-scenario combination.

between environmental exposure and human health. For example, a number of commentators have suggested that the documented relationship between daily air pollution concentrations and daily mortality rates may at least partially reflect forward displacement or harvesting.

Finally, the impacts of climate change will be felt throughout the planet. This paper's approach can be applied to data from other countries to develop estimates of the health related welfare costs of climate change elsewhere. It is especially important to develop estimates for countries where current temperatures are higher than in the United States and/or people are poorer and less able to afford life-preserving adaptations like increased energy consumption. Ultimately, the development of rational climate policy requires knowledge of the health and other costs of climate change from around the world (see e.g., Robin Burgess et al. 2011).

REFERENCES

Aroonruengsawat, Anin, and Maximilian Auffhammer. 2011. "Impacts of Climate Change on Residential Electricity Consumption: Evidence From Billing Data." In *The Economics of Climate Change: Adaptations Past and Present*, ed. Gary D. Libecap and Richard H. Steckel, 311–42. Chicago: University of Chicago Press.

Ashenfelter, Orley, and Michael Greenstone. 2004. "Using Mandated Speed Limits to Measure the Value of a Statistical Life." *Journal of Political Economy*, 112(1): S226–67.

Barreca, Alan. 2009. "Climate Change, Humidity, and Mortality in the United States." Tulane University Department of Economics Working Paper 0906.

Basu, Rupa, and Jonathan M. Samet. 2002. "Relation Between Elevated Ambient Temperature and Mortality: A Review of the Epidemiologic Evidence." *Epidemiologic Reviews*, 24(2): 190–202.

Bosello, Francesco, Roberto Roson, and Richard S. J. Tol. 2006. "Economy-Wide Estimates of the Implications of Climate Change: Human Health." *Ecological Economics*, 58(3): 579–91.

British Atmospheric Data Centre. 2007. "Hadley Climate Model 3." http://badc.nerc.ac.uk/data/.

Burgess, Robin, Olivier Deschênes, David Donaldson, and Michael Greenstone. 2011. "Weather and Death in India: Mechanisms and Implications for Climate Change." Unpublished.

Centers for Disease Control and Prevention. 1968–2002. "Compressed Mortality Files." http://www.cdc.gov/nchs/data_access/cmf.htm.

Chay, Kenneth Y., and Michael Greenstone. 2003. "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession." *Quarterly Journal of Economics*, 118(3): 1121–67.

Costa, Dora L., and Matthew E. Kahn. 2004. "Changes in the Value of Life, 1940–1980." *Journal of Risk and Uncertainty*, 29(2): 159–80.

Deschênes, Olivier, and Michael Greenstone. 2008. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the U.S." MIT Working Paper 07-19.

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US: Dataset." *American Economic Journal: Applied Economics*. http://www.aeaweb.org/articles.php?doi=10.1257/app.3.4.152.

Deschênes, Olivier, and Enrico Moretti. 2009. "Extreme Weather Events, Mortality, and Migration." *Review of Economics and Statistics*, 91(4): 659–81.

Emanuel, Kerry. 2005. "Increasing Destructiveness of Tropical Cyclones over the Past 30 Years." *Nature*, 436(7051): 686–88.

Engle, Robert F., C. W. J. Granger, John Rice, and Andrew Weiss. 1986. "Semiparametric Estimates of the Relation between Weather and Electricity Sales." *Journal of the American Statistical Association*, 81(394): 310–20.

Franco, Guido, and Alan H. Sanstad. 2006. "Climate Change and Electricity Demand in California." *Climatic Change*, 87(Suppl 1): S139–51.

Grossman, Michael. 2000. "The Human Capital Model." In *Handbook of Health Economics*, Vol. 1A, ed. Anthony J. Culyer and Joseph P. Newhouse, 347–408. Amsterdam: Elsevier Science, North-Holland.

Harrington, Winston, and Paul R. Portney. 1987. "Valuing the Benefits of Health and Safety Regulation." *Journal of Urban Economics*, 22(1): 101–12.

Hayhoe, Katharine, Daniel Cayan, Christopher B. Field, Peter C. Frumhoff, Edwin P. Maurer, Norman L. Miller, Susanne C. Moser, et al. 2004. "Emissions Pathways, Climate Change, and Impacts on California." *Proceedings of the National Academy of Sciences*, 101(34): 12422–27.

Huynen, Maud M. T. E., Pim Martens, Dieneke Schram, Matty P. Weijenberg, and Anton E. Kunst. 2001. "The Impact of Heat Waves and Cold Spells on Mortality Rates in the Dutch Population." *Environmental Health Perspectives*, 109 (5): 463–70.

Intergovernmental Panel on Climate Change (**IPCC**)**.** 2007. "Climate Change 2007: Synthesis Report." http://www.ipcc.ch/publications_and_data/publications_ipcc_fourth_assessment_report_synthesis_report.htm. (accessed May 2009).

Karl, Thomas R., and Kevin E. Trenberth. 2003. "Modern Global Climate Change." *Science*, 302(5651): 1719–23.

Klinenberg, Eric. 2002. *Heat Wave: A Social Autopsy of Disaster in Chicago*. Chicago: University of Chicago Press.

Knobel, Robin, and Diane Holditch-Davis. 2007. "Thermoregulation and Heat Loss Prevention After Birth and During Neonatal Intensive-Care Unit Stabilization of Extremely Low-Birthweight Infants." *Journal of Obstetric, Gynecologic, & Neonatal Nursing*, 36(3): 280–87.

Kunst, Anton E., Feikje Groenhof, and Johan P. Mackenbach. 1994. "The Association between Two Windchill Indices and Daily Mortality Variation in The Netherlands." *American Journal of Public Health*, 84(11): 1738–42.

Martens, W. J. M. 1998. "Climate Change, Thermal Stress and Mortality Changes." *Social Science & Medicine*, 46(3): 331–44.

National Center for Atmospheric Research. 2007. Community Earth System Model. http://www.cgd.ucar.edu/csm/. (accessed May 2007).

National Climatic Data Center. 1968–2002. "Surface Land Daily Cooperative Summary of the Day." http://cdo.ncdc.noaa.gov/pls/plcimprod.

National Institute of Environmental Health Sciences. 2010. "A Human Health Perspective on Climate Change."

Portier, Christopher J., Kimberly Thigpen Tart, Sarah R. Carter, Caroline H. Dilworth, Anne E. Grambsch, Julia Gohlke, Jeremy Hess, et al. 2010. "A Human Health Perspective On Climate Change: A Report Outlining the Research Needs on the Human Health Effects of Climate Change." The Interagency Working Group on Climate Change and Health. Research Triangle Park, NC, April. http://www.niehs.nih.gov/health/docs/climatereport2010.pdf.

Rubin, Donald B. 1987. *Multiple Imputation for Nonresponse in Surveys*. New York: John Wiley & Sons.

Semenza, Jan C., Carol H. Rubin, Kenneth H. Falter, Joel D. Selanikio, W. Dana Flanders, Holly L. Howe, and John L. Wilhelm. 1996. "Heat-Related Deaths During the July 1995 Heat Wave in Chicago." *New England Journal of Medicine*, 335(2): 84–90.

Stern, Nicholas. 2007. *The Economics of Change: The Stern Review*. Cambridge University Press.

Thaler, Richard, and Sherwin Rosen. 1976. "The Value of Saving a Life: Evidence from the Labor Market." In *Household Production and Consumption*, ed. Nestor E. Terleckyj, 265–98. Cambridge, MA: National Bureau of Economic Research.

Tol, Richard S. J. 2002. "Estimates of the Damage Costs of Climate Change. Part 1: Benchmark Estimates." *Environmental and Resource Economics*, 21(1): 47–73.

U.S. Energy Information Administration. 2006. "Monthly Electric Sales and Revenue Report with State Distributions Report (Form EIA-826)." United States Department of Energy. http://www.eia.doe.gov/cneaf/electricity/page/eia826.html (accessed May 2006).

Whitman, S., G. Good, E. R. Donoghue, N. Benbow, W. Shou, and S. Mou. 1997. "Mortality in Chicago Attributed to the July 1995 Heat Wave." *American Journal of Public Health*, 87(9): 1515–18.

World Health Organization (**WHO**)**.** 2003. *Climate Change and Human Health—Risks and Responses. Summary*. Geneva: WHO. http://www.who.int/globalchange/publications/cchhsummary/en/.